The subject matter of the invention is:

1.　A method of classifying packet delivery according to a hierarchy of service classifications, comprising the steps of:

assigning a first set of packets to a first service class to be transmitted across a first service class pathway comprising at least one link path in a network;

determining at an ingress point on the network whether the first set of packets assigned to the first service class can be sent at a transmission rate corresponding to a first service class subscription along the at least one link path in the network;

allocating sufficient bandwidth corresponding to the first service class subscription across the at least one link path in the network;

executing a class-based scheduling algorithm for the transmission of the first set of packets;

assigning weight provisions to the first set of packets; and

configuring routing parameters along the at least one link path in the network, according to implementation of the scheduling algorithm and the weight assignment provisions for the first set of packets.

2.　The method according to claim 1, further comprising the step of allocating buffer space to the first service class such that transmission of the first set of packets produces a near-zero loss packet delay.

3.　The method according to claim 1, further comprising the step of determining a cap on a total peak rate of the first service class on the at least one link path in the network.

26

4. The method according to claim 1, where the class-based scheduling algorithm for the transmission of the first set of packets comprises a weighted fair queuing scheme.

5. The method according to claim 1, where the class-based scheduling algorithm for the transmission of the first set of packets comprises a weighted round robin scheme.

6. The method according to claim 1, further comprising the step of randomly dropping packets belonging to the first service class according to a weighted random early discard scheme when the network is congested.

7. The method according to claim 1, further comprising the step of assigning a second set of packets to a second service class to be transmitted across a second service class pathway in the network.

8. The method according to claim 7, further comprising the step of determining at the ingress point on the network whether the second set of packets assigned to the second service class can be sent at a transmission rate corresponding to the second service class along each of the at least one link path in the network.

9. The method according to claim 8, further comprising the step of allocating sufficient bandwidth corresponding to the second service class across each of the at least one link path in the network.

10.   The method according to claim 7, further comprising the step of executing a class-based scheduling algorithm for the transmission of the second set of packets.

11.   The method according to claim 10, further comprising the step of assigning weight provisions to the second set of packets.

12.   The method according to claim 11, further comprising the step of configuring the routing parameters according to implementation of the scheduling algorithm and the weight assignment provisions for the second set of packets.

13.   The method according to claim 7, further comprising the step of randomly dropping packets belonging to the second service class according to a weighted random early discard scheme during congestion of the network.

14.   The method according to claim 1, where the first service class comprises a point-to-point, unidirectional service.

15.   The method according to claim 1, where the first service class comprises a point-to-point, bi-directional service.

16.   The method according to claim 15, where the point-to-point, bi-directional service comprises asymmetric bandwidths in each direction.

17.   The method according to claim 1, where the first service class comprises a point-to-multipoint, unidirectional service with a fixed aggregate rate of packet transmission.

18.   The method according to claim 1, where the first service class comprises a multipoint-to-multipoint, bi-directional service.

19.   A method of classifying packet delivery according to a hierarchy of service classifications, comprising the steps of:

assigning a first set of packets to a first service class to be transmitted across a first service class pathway comprising at least one first link path in a network;

assigning a second set of packets to a second service class to be transmitted across a second service class pathway comprising at least one second link path in a network;

assigning a third set of packets according to a third service class to be transmitted across a third service class pathway comprising at least one third link path in a network;

determining at an ingress point on the network whether the first set of packets assigned to the first service class can be sent at a transmission rate corresponding to the first service class along the at least one first link path in the network;

allocating sufficient bandwidth corresponding to the first service class across the at least one first link path in the network;

determining at an ingress point on the network whether the second set of packets assigned to the second service class can be sent at a transmission rate corresponding to the service class along each of the at least one second link paths in the network;

29

allocating sufficient bandwidth corresponding to the second service class across each of

the at least one second link paths in the network;

executing a class-based scheduling algorithm for the transmission of the first and the

second sets of packets;

assigning weight provisions to the first and the second sets of packets; and

configuring routing parameters according to implementation of the scheduling algorithm

and the weight assignment provisions for the first and the second sets of packets.

20.   The method according to claim 19, further comprising the step of determining time

average packet loss ratios for the at least one first link path, the at least one second link path and

the at least one third link path in the network for the first, the second, and the third sets of

packets.

21.   The method according to claim 19, further comprising the step of determining long time

average packet loss ratios for the first, the second and the third sets of packets.

22.   The method according to claim 19, further comprising the step of determining a

predetermined traffic mix between the first, the second, and the third sets of packets.

23.   The method according to claim 19, further comprising the step of determining a cap on a

total peak rate of the first service class on the at least one first link path of the network.

24.   The method according to claim 19, further comprising the step of allocating a buffer to store the first, the second, or the third sets of packets.

25.   The method according to claim 19, further comprising the step of determining at the ingress point on the network whether a packet corresponds to the first or the second service class by analyzing the Diffserv Codepoints (DSCP) in the first, the second or the third sets of packets' headers.

26.   The method according to claim 19, where the first service class is characterized by having a higher transmission priority than the second service class.

27.   The method according to claim 19, where the class-based scheduling algorithm for the transmission of the first set of packets comprises a weighted fair queuing scheme.

28.   The method according to claim 19, where the class-based scheduling algorithm for the transmission of the first set of packets comprises a weighted round robin scheme.

29.   The method according to claim 19, further comprising the step of randomly dropping packets belonging to the first service class according to a weighted random early discard scheme during congestion of the network.

30.   The method according to claim 19, further comprising the step of randomly dropping packets belonging to the second service class according to a weighted random early discard scheme during congestion of the network.

31.   The method according to claim 19, where the allocation of bandwidth for the first service class is accomplished by using a first-in, first-out scheme.

32.   The method according to claim 19, where the allocation of bandwidth for the second service class is accomplished by using a first-in, first-out scheme.

33.   The method according to claim 19, further comprising dropping the first set of packets at a lower probability than dropping the second set of packets.

34.   The method according to claim 19, where the first service class comprises a point-to-point, unidirectional service.

35.   The method according to claim 19, where the first service class comprises a point-to-point, bi-directional service.

36.   The method according to claim 35, where the point-to-point, bi-directional service comprises asymmetric bandwidths in each direction.

37. The method according to claim 19, where the first service class comprises a point-to-multipoint, unidirectional service with a fixed aggregate rate of packet transmission.

38. The method according to claim 19, where the first service class comprises a multipoint-to-multipoint, bi-directional service.

39. The method according to claim 19, further comprising a subscribed rate for the first service class that equals a predetermined subscribed rate.

40. The method according to claim 19, further comprising a subscribed rate for the first service class that is less than a predetermined subscribed rate.

41. The method according to claim 19, where the second service class comprises a point-to-any-point, unidirectional service.

42. The method according to claim 19, where the second service class comprises a point-to-any-point, bi-directional service.

43. The method according to claim 19, further comprising a subscribed rate for the second service class that equals a predetermined subscribed rate.

44. The method according to claim 19, further comprising a subscribed rate for the second service class that is less than a predetermined subscribed rate.

45.   The method according to claim 22, where the predetermined traffic mix between the first and the second sets of packets allows unused bandwidth allocated to the first set of packets to be used by the second set of packets after the bandwidth for the first set of packets has been allocated.

46.   The method according to claim 19, where the allocation of sufficient bandwidth for the first set of packets is such that transmission performance for the second set of packets is not deteriorated by a predetermined level.

47.   An apparatus for classifying packet delivery according to a hierarchy of service classifications, comprising:

a first system having a first processor, a first memory, a first communication connection, and a first controller connected at an ingress point of a communications network linked to a second system having a second processor, a second memory, a second communication connection, and a second controller at an egress point on the communications network by at least one link path established between a plurality of nodes and capable of transmitting packets from the ingress point, across the plurality of nodes, to the egress point on the communications network;

10   a transmitter within the first system capable of reading and assigning a first set of packets to a first service class to be transmitted across the at least one link path on the communications network;

34

the controller within the first system capable of determining at the ingress point on the communications network whether the first set of packets assigned to the first service class can be sent at a transmission rate corresponding to the first service class along each of the at least one link paths on the communications network;

a bandwidth allocation device within the first system capable of allocating sufficient bandwidth to the first service class across each of the at least one link paths on the communication network and capable of implementing a class-based scheduling algorithm executable on the controller for scheduling the transmission of the first set of packets across the communication network;

a weight provisioning mechanism within the first system capable of assigning priority classification weights to the first set of packets; and

wherein the transmitter is capable of configuring routing parameters along the at least one link path according to implementation of the class-based scheduling algorithm and the priority classification weights established by the weight provisioning mechanism for the first set of packets.

48.   The apparatus according to claim 47, where the first system is capable of executing the class- based scheduling algorithm for the transmission of the first set of packets comprising a weighted fair queuing scheme.

49.   The apparatus according to claim 47, where the first system is capable of executing the class- based scheduling algorithm for the transmission of the first set of packets comprising a weighted round robin scheme.

50.   The apparatus according to claim 47, wherein the first system is capable of randomly dropping packets belonging to the first service class according to a weighted random early discard scheme when the network is congested.

51.   The apparatus according to claim 47, wherein the first system is capable of assigning the second set of packets to a second service class to be transmitted across a second service class pathway comprising the at least one link path connecting the plurality of nodes in the network.

52.   The apparatus according to claim 51, wherein the first system is capable of determining at the ingress point on the network whether the second set of packets assigned to the second service class can be sent at a transmission rate corresponding to the second service class along each of the at least one link path connected the plurality of nodes in the network.

53.   The apparatus according to claim 51, wherein the first system is capable of executing a class based scheduling algorithm for the transmission of the second set of packets.

54.   The apparatus according to claim 53, wherein the first system is capable of assigning weight provisions to the second set of packets.

55.   The apparatus according to claim 54, wherein the first system is capable of configuring the routing parameters according to implementation of the scheduling algorithm and the weight assignment provisions for the second set of packets.

56. The apparatus according to claim 51, wherein the first system is capable of randomly dropping packets belonging to the second service class according to a weighted random early discard scheme during congestion of the network.